

CHAPTER 1

# The Trader's Edge

*If you would be a real seeker after truth, it is necessary that at least once in your life you doubt, as far as possible, all things.*

—René Descartes

There is something fascinating and mesmerizing about price movements in actively traded markets; academics, researchers, traders, and analysts are drawn to study markets, perhaps captivated as much by the patterns in the market as by the promise of financial gain. Many people believe that price changes are random and unpredictable; if this were true, the only logical course of action would be to avoid trading and to invest in index funds. This is, in fact, what a significant number of financial advisers recommend their clients do. On the other hand, there are analysts and traders who believe that they have some edge over the market, that there is some predictability in prices. This camp divides into two groups that historically have been diametrically opposed: those who make decisions based on *fundamental* factors and those who rely on *technical* factors. Fundamental analysts and traders make decisions based on their assessment of value, through an analysis of a number of factors such as financial statements, economic conditions, and an understanding of supply/demand factors. Technical traders and analysts make decisions based on information contained in past price changes themselves.

Our work here concerns the latter approach. Few traders make decisions in a vacuum; technical traders may consider fundamental factors, and fundamental traders may find that their entries and exits into markets can be better timed with an understanding of the relevant elements of market structure, money flows, and price action. Most traders find success with a hybrid approach that incorporates elements from many disciplines, and there are very few purely technical or fundamental decision makers. The key distinction, for us, is that technically motivated traders acknowledge the primacy of price itself. They know that price represents the end product of the analysis and decision

making of all market participants, and believe that a careful analysis of price movements can sometimes reveal areas of market imbalance that can offer opportunities for superior risk-adjusted profits. Building the tools for that analysis and learning how to apply them is the purpose of this book.

## DEFINING A TRADING EDGE

---

Most of the time, markets are efficient, meaning that all available information is reflected in asset prices, and that price is a fair reflection of value. Most of the time, prices fluctuate in a more or less random fashion. Though a trader may make some profitable trades in this type of environment purely due to random chance, it is simply not possible to profit in the long run; nothing the trader can do will have a positive effect on the bottom line as long as randomness dominates price changes. In theory, in a true zero-expectancy game, it should be possible to trade in a random environment and to break even, but reality is different. Trading accounts in the real world suffer under the constant drag of a number of trading frictions, transaction costs, errors, and other risks. Together, these create a high hurdle that must be overcome in order to break even. It is even possible for a trader to work with a positive expectancy system and still lose a significant amount of money to the vig.

Newer traders especially are often drawn to focus on elements of performance psychology and positive thinking. There is an entire industry that caters to struggling traders, holding out hope that if they could just get their psychological issues resolved, money would flow into their trading accounts. However, this fails to address the core problem, which is that most traders are doing things in the market that do not work. Excellent execution, risk management, discipline, and proper psychology are all important elements of a good trading plan, but it is all futile if the trading system does not have a positive expectancy. These are essential tools through which a trading edge can be applied to the market, and without which a trader is unlikely to succeed in the long run. However, none of these is a trading edge in itself.

A *positive expectancy* results when the trader successfully identifies those moments where markets are slightly less random than usual, and places trades that are aligned with the slight statistical edges present in those areas. Some traders are drawn to focus on high-probability (high win rate) trading, while others focus on finding trades that have excellent reward/risk profiles. Neither of these approaches is better than the other; what matters is how these two factors of probability and reward/risk ratio interact. For instance, it is possible to be consistently profitable with a strategy that risks many times more than what is made, as long as the win rate is high enough, or with a much lower percentage of winning trades if the reward/risk ratio compensates. In all cases, the trading problem reduces to a matter of identifying when a statistical edge is present in the market, acting accordingly, and avoiding market environments that are more random. To do this well, it is essential to have a good understanding of how markets move and also some of the math behind expectancy and probability theory.

## Expected Value

*Expected value* (or *expectancy*) is a term from probability theory that every good trader and gambler understands intuitively. For our purposes, we need to define a number of scenarios that each have a precisely defined payout (or loss), and we also need to be able to quantify the probabilities of each scenario occurring. If we are analyzing actual trading records, this can be as simple as calculating summary statistics for historical trades, but the problem is much more complicated on a look-forward basis because we have to make assumptions about how closely future conditions are likely to resemble history. Furthermore, we also need to make sure that our calculations include every possible outcome so that the probabilities sum to 1.0; this is sometimes difficult in real-world applications where unforeseeable outlier events may lurk in the future. Leaving these practical considerations aside for a moment and focusing on the underlying math, multiplying the payout of each scenario by the probability of each scenario occurring creates a probability-weighted average of the payouts, which is also called the expected value.

### The Expected Value Formula

Formally, for  $k$  possible scenarios, each with a payoff of  $x$  and associated probability  $p$ , the expected value  $E()$  is defined as:

$$E(X) = x_1 p_1 + x_2 p_2 + \cdots + x_k p_k$$

or, in alternate notation:

$$E(X) = \sum_{i=1}^k x_i p_i$$

Consider a simplified example where a trader can either make or lose 1 point with 50 percent probability of either outcome. In this example, the relevant math is:  $E(X) = .5(1) + .5(-1) = 0$ . It is important to understand precisely what expectancy tells us, which, in the case of a simplified trading or game of chance scenario, is the average amount we should win or lose on each trial. Furthermore, and this is very important, like many things in the field of probability, expectancy is valid only over a fairly large sample size. Even though our trader was playing a zero expectancy game, it is entirely possible that the trader could have had many wins or losses in a row, and could actually have accumulated a significant gain or loss at some point. In fact, it is very likely this *will* happen because random data tends to have many more strings of runs than most people would expect. Over a larger sample, it is likely that the actual value realized will begin to converge on the theoretical expected value, but distortions can and do occur.

The bottom line is that you must have an edge. If you are not trading with a statistical advantage over the market, everything else is futile. Nothing will help. Discipline, money management, execution skills, and positive thinking add great value in support

of an actual edge, but they are not edges in themselves. From a statistical standpoint, the definition of an edge is simple: can you properly identify entry and exit points in the market so that, over a large sample size, the sum of the profit and loss (P&L) from your winning trades is greater than the sum of your losing trades? The question then becomes: how do you find, develop, refine, and maintain an edge? There are many answers to that question; this book shows one possible path.

## Where Does the Edge Come From?

Many of the buying and selling decisions in the market are made by humans, either as individuals, in groups (as in an investment committee making a decision), or through extension (as in the case of execution algorithms or “algos”). One of the assumptions of academic finance is that people make rational decisions in their own best interests, after carefully calculating the potential gains and losses associated with all possible scenarios. This may be true at times, but not always. The market does not simply react to new information flow; it reacts to that information as it is processed through the lens of human emotion. People make emotional decisions about market situations, and sometimes they make mistakes. Information may be overweighted or underweighted in analysis, and everyone, even large institutions, deals with the emotions of fear, greed, hope, and regret.

In an idealized, mathematical random walk world, price would have no memory of where it has been in the past; but in the real world, prices are determined by traders making buy and sell decisions at specific times and prices. When markets revisit these specific prices, the market *does* have a memory, and we frequently see nonrandom action on these retests of important price levels. People remember the hopes, fears, and pain associated with price extremes. In addition, most large-scale buying follows a more or less predictable pattern: traders and execution algorithms alike will execute part of large orders aggressively, and then will wait to allow the market to absorb the action before resuming their executions. The more aggressive the buyers, the further they will lift offers and the less they will wait between spurts of buying. This type of action, and the memory of other traders around previous inflections, creates slight but predictable tendencies in prices.

There is no mystical, magical process at work here or at any other time in the market. Buying and selling pressure moves prices—only this, and nothing more. If someone really wants to buy and to buy quickly, the market will respond to the buying and sellers will raise their offers as they realize they can get a better (higher) price. Similarly, when large sell orders hit the market, buyers who were waiting on the bid will get out of the way because they realize that extra supply has come into the market. More urgency to sell means lower prices. More buying pressure means higher prices. The conclusion is logical and unavoidable: buying and selling pressure must, by necessity, leave patterns in the market. Our challenge is to understand how psychology can shape market structure and price action, and to find places where this buying and selling pressure creates opportunities in the form of nonrandom price action.

## The Holy Grail

This is important. In fact, it is the single most important point in technical analysis—the holy grail, if you will. *Every edge we have, as technical traders, comes from an imbalance of buying and selling pressure.* That's it, pure and simple. If we realize this and if we limit our involvement in the market to those points where there is an actual imbalance, then there is the possibility of making profits. We can sometimes identify these imbalances through the patterns they create in prices, and these patterns can provide actual points around which to structure and execute trades. Be clear on this point: we do not trade patterns in markets—we trade the underlying imbalances that create those patterns. There is no holy grail in trading, but this knowledge comes close. To understand why this is so important, it is necessary to first understand what would happen if we tried to trade in a world where price action was purely random.

## **FINDING AND DEVELOPING YOUR EDGE**

The process of developing and refining your edge in the market is exactly that: an ongoing process. This is not something you do one time; it is an iterative process that begins with ideas, progressing to distilling those ideas to actionable trading systems, and then monitoring the results. Midcourse corrections are to be expected, and dramatic retooling, especially at the beginning, is common. It is necessary to monitor ongoing performance as markets evolve, and some edges will decay over time. To be successful as an individual discretionary trader means committing to this process. Trading success, for the discretionary trader, is a dynamic state that will fluctuate in response to a multitude of factors.

## **Why Small Traders Can Make Money**

This is an obvious issue, but one that is often ignored. The argument of many academics is that you can't make money trading; your best bet is to put your money in a diversified fund and reap the baseline drift compounded over many years. (For most investors, this is not a bad plan for at least a portion of their portfolios.) Even large, professionally managed funds have a very difficult time beating the market, so why should you be able to do so, sitting at home or in your office without any competitive or informational advantage? You are certainly not the best-capitalized player in the arena, and, in a field that attracts some of the best and brightest minds in the world, you are unlikely to be the smartest. You also will not win by sheer force of will and determination. Even if you work harder than nearly anyone else, a well-capitalized firm could hire 20 of you and *that* is what you are competing against. What room is there for the small, individual trader to make profits in the market?

The answer, I think, is simple but profound: you can make money because you are not playing the same game as these other players. One reason the very large funds have

trouble beating the market is that they are so large that they *are* the market. Many of these firms are happy to scrape out a few incremental basis points on a relative basis, and they do so through a number of specialized strategies. This is probably not how you intend to trade. You probably cannot compete with large institutions on fundamental work. You probably cannot compete with HFTs and automated trading programs on speed, nor can you compete with the quant firms that hire armies of PhDs to scour every conceivable relationship between markets.

This is all true, but you also do not have the same restrictions that many of these firms do: you are not mandated to have any specific exposures. In most markets, you will likely experience few, if any, liquidity or size issues; your orders will have a minimal (but still very real) impact on prices. Most small traders can be opportunistic. If you have the skills, you can move freely among currencies, equities, futures, and options, using outright or spread strategies as appropriate. Few institutional investors enjoy these freedoms. Last, and perhaps most significantly, you are free to target a time frame that is not interesting to many institutions and not accessible to some.

One solution is to focus on the three-day to two-week swings, as many swing traders do. First, this steps up out of the noise created by the HFTs and algos. Many large firms, particularly those that make decisions on fundamental criteria, avoid short time frames altogether. They may enter and exit positions over multiple days or weeks; your profits and losses over a few days are inconsequential to them. Rather than compete directly, play a different game and target a different time frame. As Sun Tzu wrote in the *Art of War*: “Tactics are like unto water; for water in its natural state runs away from high places and hastens downward . . . avoid what is strong and strike at what is weak.”

## **GENERAL PRINCIPLES OF CHART READING**

Charts are powerful tools for traders, but it is important to think deeply about what a chart is and what it represents. Though it is possible to trade by focusing on simple chart patterns, this approach also misses much of the richness and depth of analysis that are available to a skilled chart reader. Top-level trading combines traditional left brain skills of logic, math, and analytical thinking with the intuitive, inductive skills of right brain thinking. Charts speak directly to the right brain, whose native language is pictures and images. Part of your edge as a discretionary trader comes from integrating these two halves of your being; charts are a powerful tool that can facilitate this integration and foster the growth of intuition.

Modern software packages are a mixed blessing for traders. On one hand, they have greatly increased the scope and breadth of our vision. It is not unusual for a modern trader to examine 400 or 500 charts in the course of a trading day, sometimes more than once, quickly assessing the character of a market or a set of related markets. This would not have been possible in the precomputer era, when charts had to be laboriously drawn and updated by hand. However, charting software also encourages some potentially harmful habits. It is so easy to add various plots and indicators to charts and

to tweak and change settings and time frames that some traders are forever experimenting and searching for the holy grail of technical indicators. Other traders bury price bars behind a barrage of moving averages and other indicators, thinking that complexity will lead to better trading results. Simplicity is often better than complexity. A chart is nothing more than a tool to display market data in a structured format. Once traders learn to read the message of the market, they can understand the psychological tone and the balance of buying and selling pressure at any point.

When it comes to chart setup, there is no one right way, but I will share my approach. Everything I do comes from an emphasis on clarity and consistency. Clean charts put the focus where it belongs: on the price bars and the developing market structure. Tools that highlight and emphasize the underlying market's structure are good; anything that detracts from that focus is bad. When you see a chart, you want the price bars (or candles) to be the first and most important thing your eye is drawn to; any calculated measure is only a supplement or an enhancement. Consistency is also very important, for two separate reasons. First, consistency reduces the time required to orient between charts. It is not unusual for me to scan 500 charts in a single sitting, and I can effectively do this by spending a little over a second on each chart. This is possible only because every one of my charts has the same layout and I can instantly orient and drill down to the relevant information. Consistency is also especially important for the developing trader because part of the learning process is training your eye to process data a certain way. If you are forever switching formats, this learning curve becomes much longer and steeper, and the development of intuition will be stymied. Keep the same format between all markets and time frames, and keep the setup of all of your charts as consistent as possible.

### Chart Scaling: Linear versus Log

The one exception to the principle of keeping charts consistent might be in the case of very long-term charts spanning multiple years, or shorter-term charts in which an asset has greatly increased in value (by over 100 percent). In these cases, the vertical axis of the chart should be scaled logarithmically (called "semi-log" in some charting packages) to better reflect the growth rate of the market. The idea behind a *log scale chart* is that the same vertical distance always represents the same percentage growth regardless of location on the axis.

On a very long-term chart, *linearly scaled charts* will often make price changes at lower price levels so small that they disappear and they are completely dwarfed by price changes that happened at higher levels. The linear scale also magnifies the importance of those higher-level price changes, making them seem more violent and significant than they actually were. Compare Figure 1.1 and Figure 1.2, two charts of the long-term history of the Dow Jones Industrial Average (DJIA), especially noticing the differences between the two charts at the beginning and end of the series. They seem to tell completely different stories. The first chart shows a flat and uninteresting beginning followed by violent swings

(Continued)



**FIGURE 1.1** Nothing Seems to Matter Before 1985: DJIA on a Linear Scale



**FIGURE 1.2** Investors' Actual Experiences: DJIA on a Log Scale

near the right edge of the chart, while the second, the log scale chart, shows more consistent swings throughout. Over this long history, the log scale chart is a much more accurate representation of what market participants would have experienced at any point on the chart. Remember, as a rule of thumb, there are two times when log scale charts should be used: any time you have greater than a 100 percent price increase on a chart, and for any chart showing more than two years of data, whether on daily, weekly, or monthly time frames.

## Choosing Time Frames

Discretionary traders must clearly choose and define the time frame within which they will trade, and this choice of time frames is tied into deeper questions of personality and trading style. Most of the trading ideas and principles we examine in this book can be applied to all markets and all time frames, with some adjustments, but most traders will probably find themselves best suited to a specific set of markets and time frames. Traders switching time frames or asset classes will usually undergo a painful adjustment period while they figure out how to apply their tools in the new context. For now, let's leave these important considerations behind and focus on only the mechanical issues of setting up charts to cover multiple time frames. In the end, your charts must be a tool that serves your trading style, not the other way around.

Many authors have written about the advantages of combining multiple time frames. Multiple time frames can provide context for and inform patterns on a single time frame; skilled use of multiple time frames allows traders to better manage risk and to increase the expectancy of their trading plans. Nearly all technical traders consider action and structure in other time frames, though they do this in a variety of ways. Some traders are able to infer this information from a single chart, while many others prefer to actually look at multiple charts of the same market with each chart showing a different time frame. In a scheme like this, the primary time frame of focus is called the *trading time frame (TTF)*. A *higher time frame (HTF)* chart provides a bigger-picture perspective, while a *lower time frame (LTF)* chart is usually used to find precise entry points. Other variations, with up to five or six charts, are possible, and there are many traders who use only a pair of charts. Last, though the term *time frame* seems to imply that the x-axis of the chart will be a time scale (minutes, hours, days, etc.), the same proportional relationships can be applied to tick, volume, or any other activity-based axis scale on the x-axis.

In general, time frames should be related to each other by a factor of 3 to 5. There is no magic in these ratios, but the idea is that each time frame should provide new information without loss of resolution or unnecessary repetition. For instance, if a trader is watching a 30-minute chart, a 5- or 10-minute chart probably provides new information about what is going on inside each 30-minute bar, whereas a 1-minute chart would omit significant information. Using a 20-minute chart in conjunction with a 30-minute chart probably adds no new information, as the two charts will be very similar. One lesser-known relationship is that all vertical distances on charts scale with the square

root of the ratio of the time frames. This has implications for risk management, profit targets, stops, and volatility on each time frame. For instance, if a trader has been trading a system on 5-minute charts with \$0.25 stops and wishes to transfer that to 30-minute charts, the stops will probably need to be adjusted to about \$0.61 ( $\$0.25 \times \sqrt{30/5}$ ). This relationship does not hold exactly in all markets and all time frames, but it is a good rule of thumb and can give some insight into the risks and rewards of other time frames.

The rule of consistency also applies to choice of time frames. Once you have settled on a trading style and time frame, be slow to modify it unless you have evidence that it is not working. This story will be told with the most clarity and power in a consistent time frame. In addition, if you catch yourself wanting to look at a time frame you never look at while you are in a losing trade, be very careful. This is often a warning of an impending break of discipline.

### **Bars, Candles, or Other Choices**

Most traders today seem to be focused on using candlestick charts, but the more old-fashioned bar charts should not be overlooked. Both chart types display the same data points but in a slightly different format; they have the same information on them, so one is not better than the other. The main advantage of bar charts is that they can be cleaner visually and it is usually possible to fit more data in the same space because bars are thinner than candles. For many traders, the colors of candlestick charts make it easier to see the buying and selling pressure in the market, providing another important visual cue that helps the trader process the data faster.

Another issue to consider, particularly with intraday charts, is how much importance should be attached to the closing print of each period. Historically, this was *the* price in many markets, and it still has significance in some contexts. Profits and losses (P&Ls), margins, and various spreads are calculated off daily settlement prices; exchanges have complex procedures for calculating these prices, which are rarely simply the last print of the session. However, times are changing. In currencies, most domestic platforms report a closing price sometime in the New York afternoon, and we have to wonder just how important that price is for the Australian dollar or the yen, whose primary sessions ended many hours earlier. As more and more markets go to 24-hour sessions, the importance of this daily settlement price will continue to decline. The problem is even more significant on intraday bars, bar as closing prices on intraday bars are essentially random samples and may differ from platform to platform. If you are trading candlestick patterns, which attach great significance to the close, you are trading the patterns you see on your screen. If you switched to a different data provider, the data might be time-stamped differently, and you would see different patterns. How important can those patterns really be?

## **INDICATORS**

---

Indicators are calculated measures that are plotted on price charts, either on top of the price bars or in panels above or below the bars. There are many different indicators in